

Multipath RDMA Instrumentation

Brett McMillian, Dennis Ferguson, and Gary McMillian, Ph.D.

Crossfield Technology, Austin, Texas

Donald R. Snyder, III

Air Force Research Laboratory, Eglin AFB, Florida

Multipath remote direct memory access (RDMA) is a new networking technology that provides an order-of-magnitude increase in bandwidth through a full-mesh backplane and through a standard Ethernet switch fabric. This article first describes full-mesh backplane technology, RDMA access over internet protocol, and Multipath RDMA, then introduces two novel Multipath RDMA instrumentation systems: an instrumentation module and an instrumentation blade. The instrumentation module supports Giga-sample per second (GSPS) instrumentation and operates at the end of the wire using Power-over-Ethernet Plus. The instrumentation blade goes directly into a Multipath RDMA-enabled blade chassis, supports multi-GSPS instrumentation, and gives rise to real-time or near-real-time parallel processing of data produced by multi-GSPS instrumentation.

Key words: Data acquisition; Ethernet; instrumentation; hardware-in-the-loop; Multipath RDMA; power over Ethernet; remote direct memory access (RDMA).

Novel networking, instrumentation, and computation hardware is introduced for multiple Giga-sample per second (GSPS) data acquisition, hardware-in-the loop simulation, signal processing, and spectrum analysis. The system builds on 10 Gbits/s remote direct memory access (RDMA) over internet protocol (IP) and uses full-mesh backplanes to create a new networking technology called Multipath RDMA. A Multipath RDMA full-mesh bridge chip provides over an order-of-magnitude increase in bandwidth between computing nodes and supports real-time or near-real-time parallel processing in multi-GSPS instrumentation systems.

Full-mesh networking

A full-mesh bridge chip enables Multipath RDMA through a full-mesh passive backplane. The backplane incorporates a full-mesh network topology into a passive backplane, or midplane, which provides every processor blade a direct connection to every other processor blade. *Figure 1a* illustrates this full-mesh topology while *Figure 1b* shows how the same blades are also mapped to a dual-redundant star topology, the network topology used in a typical Ethernet switch fabric. Each switch provides a cut through switch connection between any two blades in the system. The full-mesh fabric on the

other hand allows for a blade to have as many connections as there are other blades in the backplane. Hence, the full-mesh fabric provides much higher aggregate bandwidth into and out of any one node.

Both the AdvancedTCA (PICMG® 2003) and VITA 46/48 standards (VITA 2008) currently define full-mesh backplanes. These backplanes support up to 16 slots with 15 serial communications channels per slot. Each channel supports full-duplex 10G Ethernet with a 10 Gigabit physical layer comprised of eight differential pairs, each operating at a 3.125 GHz signaling rate. Hence each slot has a maximum 150 Gbits/s of bandwidth in both directions through the backplane. Unfortunately, little exists on the market today that can use this bandwidth outside of the backplane. Current switch chips support the full bandwidth through the backplane but bottleneck down to one or two 10G Ethernet connections to the processors on the blades. A full-mesh bridge chip with a native host processor interface and Multipath RDMA capability is designed to alleviate this bottleneck and provide over an order-of-magnitude greater bandwidth between processors.

The full-mesh bridge chip, shown in *Figure 2*, provides a bridge between a processor's high-speed host interface and a full-mesh backplane interface. The host interface can take the form of any high-speed,

Report Documentation Page			Form Approved OMB No. 0704-0188		
Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.					
1. REPORT DATE SEP 2008		2. REPORT TYPE		3. DATES COVERED 00-00-2008 to 00-00-2008	
4. TITLE AND SUBTITLE Multipath RDMA Instrumentation				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S)				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Air Force Research Laboratory, Eglin AFB, FL 32542				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT Same as Report (SAR)	18. NUMBER OF PAGES 9	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			




Figure 1. (a) Full-mesh network topology, (b) dual-redundant star network topology

chip-to-chip interconnect technology like Advanced Micro Device's (AMD's) HyperTransport™ (HT) (HyperTransport Consortium 2001), Intel's Quick-Path Interconnect (QPI) (Intel 2008), and Rambus' FlexIO™ (Rambus 2008), which is currently used in the IBM, Sony, and Toshiba Cell Broadband Engine (Cell), or other future interconnect standard. Dual host interfaces support dual-processor blades, with the full-mesh bridge providing a peak bandwidth of around 20 GBytes/s (160 Gbits/s) to each processor. The backplane interface supports the maximum fifteen channels, which means that two processors connected via host interfaces will be able to share the full 150 Gbits/s through the backplane.

Remote direct memory access over IP

In addition to the fifteen serial channels to the backplane, the full-mesh bridge chip provides two interfaces for use as external 10G Ethernet ports into a dual redundant switch fabric as shown in *Figure 3*. In other words, the full-mesh bridge chip has the equivalent of 17 10G serial connections: two external Ethernet ports and 15 backplane channels. Each processor can use these 17 10G connections to transfer data directly from its memory to another processor's memory on any other blade in the system via RDMA.

RDMA over IP (RDMA Consortium 2003) differs from traditional Ethernet transmission in that RDMA eliminates unnecessary buffering in the operating

system when transmitting and receiving packets. Instead of copying packets into a buffer in the operating system before sending, an RDMA Network Interface Controller (RNIC) takes data directly from application user-space memory, applies the appropriate network layer protocols and Ethernet link layer frame and sends the packet across the network. On the receiving end, another RNIC receives the packet and places the payload directly into application user-space memory. By removing the unnecessary data copying in the kernel and off-loading network protocol processing at both ends of the link, RNICs alleviate the latency issues normally associated with Ethernet and make Ethernet a viable solution for high-speed, low-latency instrumentation systems.

To get into more detail, the RDMA Consortium (RDMA Consortium 2003) defined a suite of protocols at the transport layer that enables cooperating Direct Memory Access (DMA) engines at each end of a network connection to move data between user-space memory buffers with minimal support from the kernel and with "zero copy" to intermediate buffers. The Remote Direct Memory Access Protocol Verbs specification (Recio 2007) describes the behavior of the protocol off-load hardware and software, defines the semantics of the RDMA services and how they appear to the host software, and defines both the user and kernel application programming interface. The Verbs specification defines RDMA *READ*, *WRITE*,




Figure 2. Full-mesh bridge chip

and *SEND* operations that transport data between user space memories or into a receive queue, respectively, and defines send and receive queues and queue pairs to control data transport and completion queues to signal when an operation is complete. Work requests are converted into work queue elements which are processed in turn by the off-load engine. An asynchronous event (interrupt) signals when work is complete. Also, data need not be in contiguous locations in either the source or destination memories as scatter-gather lists define the physical memory locations of data segments.

The full-mesh bridge includes built-in support for RDMA over IP, and the processors control the device through their host interface. However, since the RNIC accesses processor main memory directly, the RNIC must have a high bandwidth path to main memory in order to not bottleneck. XDR and DDR3 are the only memory interfaces that currently match the 20 Gbps bandwidth of advanced host interfaces. The Cell Broadband Engine, which uses XDR memory, is the only processor that currently provides both a host and memory interface at this bandwidth, but Intel and AMD should reach this level of performance in the near term.

Multipath RDMA

While full-mesh connectivity provides high aggregate bandwidth between multiple processors distributed across the backplane, the bandwidth between any two processors is the same as if the backplane were a simple switch because only one direct connection exists between any two slots in the backplane. The same is true for connections between processors connected through an external dual-redundant Ethernet switch fabric. Processors connected across the dual-redundant switch fabric have at most 20 Gbits/s of bandwidth for RDMA data transfer (using both paths through the redundant fabrics). However, a new technology called Multipath RDMA, incorporated into the bridge chip, allows two individual processors located anywhere in the network to use the full 150 Gbits/s bandwidth through the backplane plus the 20 Gbits/s bandwidth through the external switch fabric for RDMA data transfer via multiple parallel paths.




Figure 3. Full-mesh bridge chip with Ethernet ports

Multipath RDMA gets its name from its similarity to multipath in wireless communication. In wireless communications, multipath refers to the multiple paths traveled by electromagnetic waves between antennas due to reflections in the environment. New Multi-Input, Multi-Output (MIMO) technologies like those incorporated into IEEE 802.11n (IEEE 2008) exploit multipath transmission to increase the overall data rate. Similarly, Multipath RDMA uses every possible route through both the full-mesh backplane and the dual-redundant switch fabric to route packets between any two processors. The most significant improvement Multipath RDMA offers over current networking technology is that it provides over an order-of-magnitude increase in bandwidth between processors in the same backplane and between blades connected through a dual-redundant switch fabric.

First, when processor blades share a common full-mesh backplane, they have 17 available channels to transmit data to any other processor. Any two processors in the network have one direct channel through the full-mesh backplane plus two switched connections. However, with Multipath RDMA the two processors also gain 14 one-hop connections through the remaining full-mesh channels via the other 14 blades' full-mesh bridge chips. The full-mesh bridge chips act as cut-through switches between the sending and receiving full-mesh bridges on the source and destination blades. *Figure 4* shows an example of Multipath RDMA between nodes in the same backplane.

Second, when two blade chassis are connected through a dual-redundant switch fabric, Multipath RDMA exploits the two full-mesh backplanes to increase the bandwidth through the switch fabric. On the sending side, the full-mesh bridge chip segments the payload and transmits the resulting packets through all 17 available channels. The two switch fabric ports send their packets to the Ethernet ports on the receive side in the standard way. The 15 connections through the backplane interface, however, route through the other 15 full-mesh bridge chips in




Figure 4. Multipath RDMA through a full-mesh backplane

the backplane, which act as cut-through switches, to the switch fabric ports on each blade. Then the packets are sent across the switch fabric simultaneously to the 15 blades connected to the receiving blade through the other backplane. The full-mesh bridge chips on those blades then act as cut-through switches to the channels connected to the receiving blade's full-mesh bridge chip. The full-mesh bridge on the receiving blade then combines the payloads from each packet and places the data directly into the receiving processor's memory. Hence for a Multipath RDMA transmission through the dual-redundant switch fabric, a processor has two one-hop switched connections and 15 three-hop connections through full-mesh bridges. *Figure 5* shows an example of Multipath RDMA between nodes in separate full-mesh backplanes.

In both cases, the bandwidth between any two full-mesh bridge chips is a multiple of 10 Gbits/s up to the maximum 170 Gbits/s. The number of channels utilized by Multipath RDMA can change on the fly such that the number of communication paths between any two processors is dynamic. As a result, Multipath RDMA can support a large number of connections at dynamically changing bandwidths in order to optimize connections for a wide range of algorithms. When compared to current commercial off-the-shelf systems, Multipath RDMA provides more than 16 times the maximum bandwidth between processors connected through a full-mesh backplane and more than eight times the maximum bandwidth between processors connected through, and utilizing, a dual-redundant switch fabric. However, real-time or near-real-time systems can only

take advantage of this bandwidth if the instruments are able to obtain the same amount of bandwidth into and out of the network fabric.

Multipath RDMA instrumentation

In order to achieve high bandwidth between instrumentation and a Multipath RDMA-enabled supercomputing cluster, a new instrumentation system had to be architected from the ground up that could incorporate the concept of using multiple connections to achieve higher bandwidth. First, the instrumentation systems must accommodate analog-to-digital converters (ADCs), digital-to-analog converters (DACs), or other sensors or actuators that source or sink high bandwidth data streams. With up to 170 Gbits/s bandwidth through a blade chassis, Multipath RDMA supports ADCs and DACs with multi-GSPS data rates. Second, the instrumentation systems must either have a single high bandwidth connection into the computing cluster or several lower bandwidth parallel connections.

To this end, designs for two different instrumentation systems developed. The first design is a 10/20 Gbits/s RDMA Power-over-Ethernet Plus (PoEP) Instrumentation Module that supports ADCs and DACs up to 2 GSPS (at 8 bits of resolution). The second design approach puts the ADCs and DACs directly onto the blades in the supercomputing cluster and supports sampling rates up to 20 GSPS (at 8 bits of resolution). The first approach has the advantage of placing the ADCs and DACs close to the sensors and actuators, which eliminates noise on the analog side and decreases the number of cables. The second

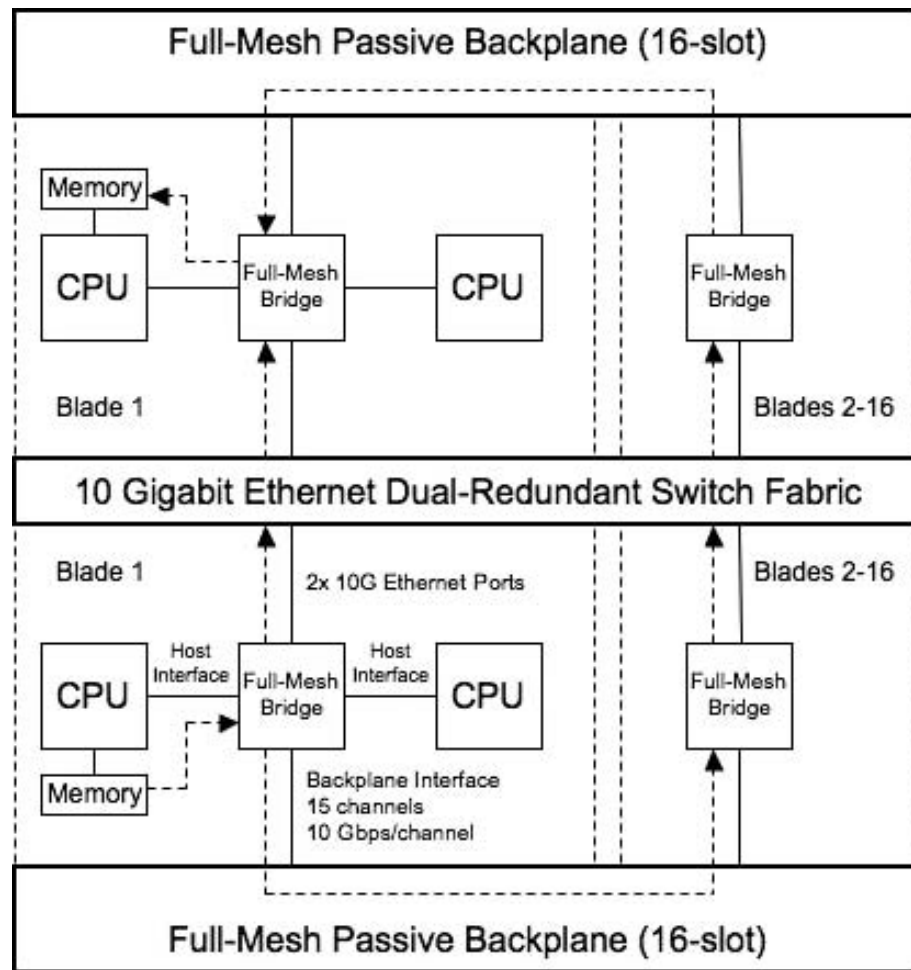


Figure 5. Multipath remote direct memory access (RDMA) across a dual-redundant switch fabric

approach has the advantage of supporting the fastest ADCs and DACs currently on the market and provides enough bandwidth into a full-mesh backplane to distribute full-rate sensor data across parallel signal processors using Multipath RDMA.

10/20 Gbits/s RDMA PoEP instrumentation modules

The 10/20 Gbits/s RDMA PoEP Instrumentation Module is a remote high-speed sensor data acquisition and actuator controller system that both transmits data through and gets its power from a single CAT6 or CAT7 Ethernet cable. The module consists of a motherboard and an XMC card, the first of which is responsible for handling the Ethernet communication, IEEE 1588 time synchronization, and main power supplies. The XMC card contains high-speed ADCs or DACs and analog instrumentation. With pluggable XMC cards, the module is extremely versatile and can be used in almost any data acquisition, signal

processing, hardware-in-the-loop, or other type of instrumentation system.

XMC is the VITA 42 standard (VITA 2008) that extends the PCI Mezzanine Card (PMC) to several different high-speed buses including PCI Express (VITA 42.3) and HyperTransport (VITA 42.4). The XMC standard supports up to 16 data lanes per slot. With 8-lane PCI Express V2.0 the XMC slot provides up to 32 Gbits/s of bandwidth, which means that a plug-in card can support a single ADC sampling at 4 GSPS with 8-bits of resolution or a 2 GSPS DAC with 16-bits of resolution or multiple ADCs or DACs with lower sampling rates or resolutions. However, the module itself only supports up to 20 Gbits/s into a switch fabric, which means the XMC card is limited to a single ADC sampling at 2 GSPS with 8 bits of resolution or a 1-GSPS DAC with 16 bits of resolution.

XMC cards with various ADC and DAC configurations are currently available as commercial off-the-shelf boards. The real innovation in the instrumenta-

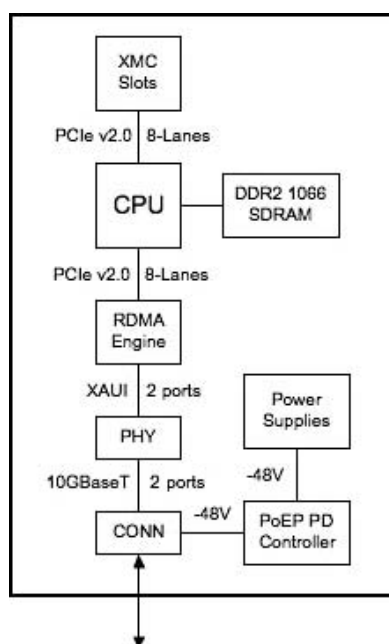


Figure 6. 10G PoEP motherboard

tion module is the motherboard. The motherboard supports up to two 10G Ethernet ports, IEEE 1588, and PoEP, and has an extremely small form factor about the same size as a double-wide XMC card (15 cm \times 15 cm), and as a result, can be placed relatively close to the sensors or actuators, which reduces the lengths of the analog sensor and actuator wires. Figure 6 shows a simplified block diagram of the motherboard.

A dual-port RDMA ASIC, like NetEffect's 10GbE ASIC (NetEffect Inc., Austin, Texas) for example, on the motherboard provides two 10 Gbits/s RDMA over IP Ethernet ports. With RDMA support, the instrumentation module can transfer data directly between its own memory and tens, hundreds, or even thousands of processor memories in a supercomputing cluster through an Ethernet switch fabric. While only two 10G Ethernet connections may not warrant using Multipath RDMA, it can be used on the server end to send data to multiple instrumentation modules from a single processor or read data from multiple instrumentation modules into a single processor.

IEEE 1588 Precision Time Protocol (PTP) (IEEE 2002) is a time synchronization protocol that allows the motherboard to synchronize a local clock with a master clock located somewhere in the network. Once the motherboard clock is synchronized with the IEEE 1588 PTP master clock, it can then be used as a reference for the ADC and DAC clocks on the XMC card. By using IEEE 1588 PTP time synchronization, every instrumentation and processing node in the

network can have clocks synchronized to within 100 ns of each other or better.

PoEP is the IEEE 802.3at (IEEE 2005) extension to the IEEE 802.3af (IEEE 2003) Power-over-Ethernet standard that allows a powered device to use up to 70 W of power supplied by power sourcing equipment. The IEEE 802.3at draft standard defines a midspan power sourcing equipment device that connects in series to insert power into an Ethernet link. Since the motherboard accepts PoEP, the motherboard and XMC card can use up to 70 W of power, which is more than enough power for most applications. PoEP products are currently available for Gigabit Ethernet, but extending PoEP technology to 10G Ethernet is straightforward since 10GBase-T uses the same 4 pairs as 1000Base-T.

As mentioned above, the instrumentation module connects to the Multipath RDMA enabled cluster and receives its power via a PoEP midspan device or switch. With a midspan configuration and CAT7 cable, each instrumentation module can be located up to 100 m from a switch. Hence, the instrumentation modules can be located large distances from the processing cluster, which is often convenient when the cluster cannot be physically located nearby. Figure 7 shows an illustration of the module network.

The only disadvantage of the instrumentation module approach is the limited bandwidth between the ADCs and DACs and the cluster. With only 20 Gbits/s maximum bandwidth into the switch fabric, the modules cannot support the sampling rate and resolution of the fastest ADCs and DACs currently available. For this reason, a second, faster instrumentation system approach was developed that has sufficient bandwidth to support state-of-the-art ADCs and DACs and can use Multipath RDMA to distribute data throughout the cluster.

Multipath RDMA instrumentation blades

In order to achieve the maximum bandwidth into a Multipath RDMA system, the ADCs and DACs must have a high-speed interface into a full-mesh backplane. A straightforward way to gain access to the full-mesh backplane is if the ADCs and DACs are on the blades themselves. The ADCs and DACs can be interfaced to the full-mesh bridge using an FPGA, which buffers data in memory and provide a host interface to the full-mesh bridge.

FPGAs can communicate with the fastest ADCs and DACs on the market, through a high speed interface. The latest FPGAs provide a DDR3 interface supporting up to several Gigabytes of memory, which, for example, can be used to buffer data from an ADC before it is distributed to processors in the cluster.

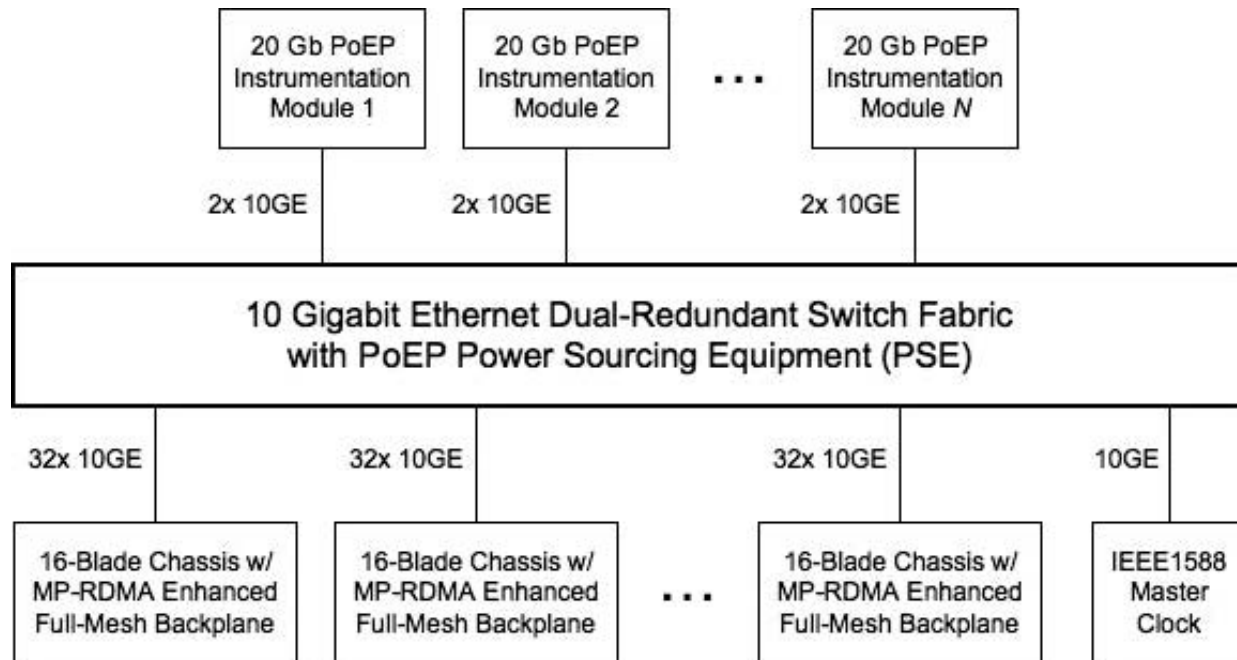


Figure 7. Instrumentation module network with the IEEE 1588 master clock

Finally, the FPGA's host interface bridges the instrumentation into the full-mesh bridge. The full-mesh bridge is then able to use Multipath RDMA to transfer data between the FPGA's memory and the memory of any processor in the cluster. Furthermore, with IEEE 1588 PTP support built into the full-mesh bridge, instrumentation blades and processing blades will all support clock time synchronization in the same fashion as the instrumentation modules. *Figure 8* shows a block diagram of the instrumentation blade architecture.

Application example

As a demonstration of how the Multipath RDMA instrumentation system works, the following section shows how the system could compute a billion-point complex fast Fourier transform (CFFT) of acquired sensor data in near-real-time using the fastest commercial off-the-shelf ADCs and multiple Cell Broadband Engines from IBM. First, assume two 8-bit 5-GSPS ADCs capture an arbitrary waveform starting at time $t = 0$ seconds on an instrumentation blade connected to a full-mesh backplane. The two ADCs sample the waveform in-phase and quadrature-phase (I & Q). A billion complex data points are sampled every 214.7 ms (a billion in this case is actually 2^{30}). This data rate corresponds to a bandwidth of 10 GBytes/s from the ADCs. Since the ADC interface and the DDR3 interface on the FPGA have more than 10 GBytes/s bandwidth, the

data points are stored in the FPGA memory as fast as they are sampled.

As the data is stored in memory, the FPGA simultaneously tasks the full-mesh bridge to transport the data to a Cell processor somewhere in the cluster using Multipath RDMA. Since the host interface and the network bandwidth to a Cell processor through the full-mesh bridge using Multipath RDMA are greater than the sampling rate of the ADC, the data is streamed into a Cell processor's memory as fast as the ADCs sample the waveform.

The IBM Cell Broadband Engine is a multicore heterogeneous processor with one PowerPC Processor Element (PPE) and eight Synergistic Processor Elements (SPEs). The most recent update to the Cell, the PowerXCell 8i (IBM, Los Alamos, New Mexico), used in the Los Alamos Roadrunner QS22 blade cluster (Koch 2007), has a high speed FlexIO interface that supports up to 20 GBytes/s I/O bandwidth and a DDR2 memory controller that supports up to 64 Gigabytes of memory. This latest update also boasts 230 GFLOPS single precision floating-point performance and over 100 GFLOPS double-precision floating point performance.

In order to perform a near-real-time billion-point CFFT on a Cell, the processor must have both enough memory to support the large data set and a significant amount of computational power to perform the calculations in the required time. A billion-point single-precision floating-point CFFT requires 8 Bytes

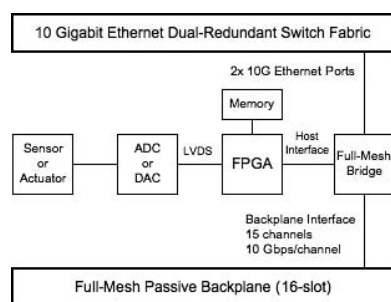


Figure 8. Multipath remote direct memory access (RDMA) instrumentation blade

of storage per complex data point and thus needs 8 GB of memory for storage. While the first-generation Cell's XDR interface did not support this much memory, the PowerXCell 8i's DDR2 controller interface can easily meet this requirement.

The Cell provides unmatched computational power in a single chip. With 230 GFLOPS single-precision performance, the Cell can perform a 16M-point CFFT in 0.043 seconds (Chow, Fossum, and Brokenshire 2005). The time required to compute a 16M-point CFFT on a Cell processor can be scaled to estimate the time required to calculate a billion-point CFFT. The first part of the calculation involves determining the relative complexity factor of a billion-point CFFT to a 16M-point CFFT. The complexity of a CFFT is determined by $N \log_2(N)$, where N is the number of points. Hence the relative complexity of a billion-point CFFT when compared to a 16M-point CFFT is $[1B \log_2(1B)]/[16M \log_2(16M)]$ or 80. Assuming the billion-point CFFT is as equally parallelizable as the 16M-point CFFT the time a Cell takes to calculate a billion-point CFFT is 80×0.043 seconds or 3.44 seconds.

Since the ADC can sample one billion complex data points in 214.7 ms, computing the billion-point CFFT in real time is not currently possible on a single Cell. However, by distributing the computational load across multiple Cell processors, the CFFT can be performed in real-time or near-real-time. A Multipath RDMA Instrumentation system with one instrumentation blade and nine Cell blades containing 18 Cells is able to compute a billion-point CFFT on a 5 GSPS complex signal in real-time.

Conclusion

In conclusion, Multipath RDMA is a new networking technology that exploits multiple routes through a fabric to achieve an order-of-magnitude increase in bandwidth between processors in instrumentation systems and supercomputing clusters. It also provides a novel means to interface ultra-high-speed instrumentation to a computing cluster and makes real-time

or near-real-time data acquisition and parallel processing possible. Multipath RDMA is a quantum leap in the battle for more bandwidth. While several organizations are currently working on 100G Ethernet standards, Multipath RDMA offers even greater bandwidth using current 10G Ethernet technology. As Ethernet speeds increase, Multipath RDMA will maintain this advantage. Multipath RDMA is an innovative way to get to the next level of bandwidth capacity by exploiting an underutilized resource in a blade system, the passive backplane. □

BRETT McMILLAN is an engineer at Crossfield Technology in Austin, Texas. He has worked on several programs in the areas of circuit design, instrumentation, high-speed networking, Linux driver development, analog and mixed-signal design, and wireless systems. He codeveloped the concept of Multipath RDMA and also recently developed a PoEP instrumentation module similar to the module described in this article for strain gauge, pressure, and temperature sensors as part of a program with the Arnold Engineering Development Center. His areas of expertise include printed circuit board design and computer networking. He holds a B.A. in physics from Grinnell College, Grinnell, Iowa. E-mail: brett.mcmillan@crossfieldtech.com

DONALD R. SNYDER is a senior scientist with the Air Force Research Laboratory Munitions Directorate and is the lead scientist for the Advanced Guidance Division Kinetic Kill Vehicle Hardware in the Loop Facility (KHILS). He currently directs a research portfolio of over 30 research efforts, including hardware-in-the-loop test technologies. He is a fellow of SPIE and a member of IEEE/Laser and Electro Optics Society, Society for Information Display, NDIA, National Training System Association, and SPIE High Speed Instrumentation Working Group. Snyder is the recipient of the ITEA Modeling and Simulation Award for MDA Modeling and Simulation Scene Generation. E-mail: snyder@eglin.af.mil

DENNIS FERGUSON has held technology management and engineering positions at several companies over his 32-year career. At Crossfield Technology, Ferguson has worked on software radios and harsh environment wireless instruments. At Honeywell, he led many technology developments involving integrated circuits and software radios for commercial aviation and communications. Ferguson holds a BSEE and an MBA. He has been a member of IEEE for 36 years and is currently vice chair of the Central Texas Solid State Circuits society. Ferguson has been awarded six patents and has published or presented over 37 articles. E-mail: dennis.ferguson@crossfieldtech.com

GARY McMILLAN is a scientist and partner at Crossfield Technology in Austin, Texas. He has served as principal

investigator for numerous R&D programs in the areas of instrumentation, high-speed networking and signal processing. His technical areas of expertise include sensors, instrumentation, systems engineering, signal processing, wireless systems, and network data communications. He is a member of the IEEE and previously served as the secretary of the Central Texas Section of the IEEE Communications/Signal Processing Society. He holds a B.S. in physics from Sam Houston State University, Huntsville, Texas and a M.A. and Ph.D. in physics from Rice University, Houston, Texas. E-mail: gary.mcmillan@crossfieldtech.com

References

Chow, A. C., Fossum, G. C. and Brokenshire, D. A. 2005. "A Programming Example: Large FFT on the Cell Broadband Engine." White Paper available at <http://www-01.ibm.com/chips/techlib/techlib.nsf/techdocs/0AA2394A505EF0FB8725>.

HyperTransport™ Consortium. 2001. Hypertransport Link Specification. Sunnyvale, California: Hypertransport Consortium. Available at <http://www.hypertransport.org>. Accessed June 6, 2008.

IEEE (Institute of Electrical and Electronics Engineers). 2002. Standard for Precise Time Synchronization as the Basis for Real Time Applications in Automation. Piscataway, New Jersey, IEEE available at <http://ieee1588.nist.gov/> or <http://www.ieee1588.com>. Accessed June 6, 2008.

IEEE (Institute of Electrical and Electronics Engineers). 2003. IEEE Std. 802.3af—IEEE Standard for Information technology—Telecommunications and information exchange between systems—Local and metropolitan area networks. Specific requirements. IEEE Computer Society.

IEEE (Institute of Electrical and Electronics Engineers). 2005. IEEE P802.3at DTE Power Enhancements Task Force authorization. <http://www.ieee802.org/3/at/>

IEEE (Institute of Electrical and Electronics Engineers). 2008. Wireless Networking Standard 802.11n, proposed amendment n, Multi-Input, Multi-Output (MIMO) technologies. New York: IEEE.

Intel. 2008. White Paper Intel QuickPath Architecture. Santa Clara, California: Intel. Available at http://www.intel.com/pressroom/archive/reference/whitepaper_QuickPath.pdf Accessed June 6, 2008.

Koch, K. 2007. Roadrunner System Overview. Powerpoint briefing. Available at [http://www.lanl.gov/orgs/hpc/roadrunner/rrinfo/RR%20webPDFs/Roadrunner%20System%20Overview%20Oct%2017,%202007%20\(LAUR\).pdf](http://www.lanl.gov/orgs/hpc/roadrunner/rrinfo/RR%20webPDFs/Roadrunner%20System%20Overview%20Oct%2017,%202007%20(LAUR).pdf). October 17, 2007.

PICMG® (PCI Industrial Computers Manufacturers Group). 2003. Advanced TCA Specifications for Next Generation Telecommunications Equipment. Wakefield, Massachusetts: PICMG. Available at <http://www.picmg.org/v2internal/newinitiative.htm>. Accessed June 6, 2008.

Rambus. 2008. FlexIO™ Processor Product Brief. Available at http://www.rambus.com/assets/documents/products/FlexIO_ProcessorBus_ProductBrief.pdf. Accessed June 6, 2008.

RDMA Consortium. 2008. Architectural Specifications for RDMA over TCP/IP. available at <http://www.rdmaconsortium.org>.

Recio, R. 2007. Remote Direct Memory Access Protocol Verbs specification. RDMA Consortium. available at <http://www.rdmaconsortium.org>.

VITA. 2008. VPX and REDI Specifications. Fountain Hills, Arizona: VITA. Available at <http://www.vita.com>. Accessed June 6, 2008.

Acknowledgments

This research was sponsored by the U.S. Department of Defense Missile Defense Agency under U.S. Army contract number W9113M-06-C-0163.